Prostate Implant Reconstruction with Discrete Tomography

Xiaofeng Liu[1], Ameet K. Jain[1,2], and Gabor Fichtinger[1]

[1] Department of Computer Science, Johns Hopkins University, Baltimore, MD, USA
[2] Philips Research North America, Briarcliff, NY, USA

Abstract. We developed a discrete tomography method for prostate implant reconstructions using only a limited number of X-ray projection images. A 3D voxel volume is reconstructed by back-projection and using distance maps generated from the projection images. The true seed locations are extracted from the voxel volume while false positive seeds are eliminated using a novel optimal geometry coverage model. The attractive feature of our method is that it does not require exact seed segmentation of the X-ray images and it yields near 100% correct reconstruction from only six images with an average reconstruction accuracy of 0.86 mm (std=0.46mm).

1 Introduction

Brachytherapy is a definitive treatment for low risk prostate cancer that represents the vast majority of new cases diagnosed nowadays. The brachytherapy procedure entails permanently implanting small radioactive seeds into the prostate. The main limitation of contemporary brachytherapy is faulty seed placement that may result in insufficient dose to the cancer and/or inadvertent radiation to the rectum, urethra, and bladder. Intra-operative implant optimization promises a major clinical breakthrough, but for this technique to succeed the implanted seeds must be reconstructed and registered with the anatomy [1]. This work concentrates on the first problem, reconstruction.

C-arm X-ray fluoroscopy is the gold standard in observing brachytherapy seeds and therefore is a natural candidate for implant localization. The 3D coordinates of the seeds can be calculated from multiple X-ray images upon solving the correspondence problem [2-6]. These methods uniformly require that seeds are accurately segmented in the X-ray images. Significant research has been dedicated to this problem, still without clinically robust and practical solution. To make the problem worse, typically 7%, but often as much as 43% of the seeds can be hidden in the X-ray images [7], and the recovery of these seeds is an exigent task that often leaves seeds undetected. Su *et al.* [7] proposed a solution to the hidden seed problem by resolving seed clusters and extending previously published approach [3], but it still required perfectly localizing all visible seeds all projection images.

Classic tomosynthesis might seem a suitable reconstruction, but unfortunately, it is impractical in brachytherapy because (1) the swing space of the C-arm is limited due to collision hazards and (2) the number of X-ray images is strictly limited in order to

save the patient and the OR crew from excessive toxic radiation. Tutar *et al.* [8] has proposed a variant of tomosynthesis denoted as selective back projection. However this method demands a large number of images ($\geq$7) and wide C-arm angle ($\geq$25°) to succeed. It is also prone to introducing false positive (FP) seeds, which from a dosimetric point of view are more troublesome than hidden seeds, because they act toward underdosing the cancer. Tutar *et al.* use a heuristic rule to recognize FP seeds by their sizes, but since C-arm pose estimation and calibration errors affect the size of objects, this may result in faulty separation of the true and false seeds.

Our approach using discrete tomography is different in that after generating a 3D volume using back-projection we detect and remove the false positive seeds by solving an optimal coverage problem. We achieved high reconstruction rate with fewer images.
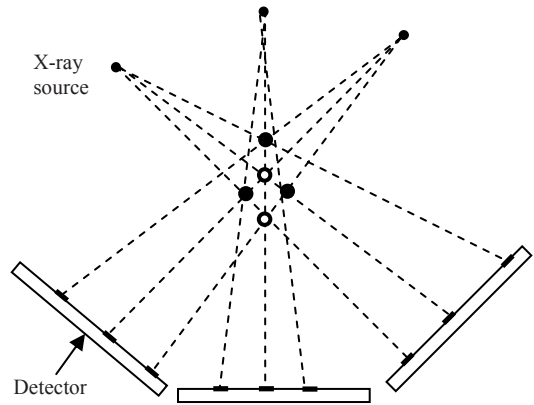
X-ray
source

Detector

Fig. 1. Two false positive seeds are introduced by three legitimate seeds when their projectiles inter-sect the same voxels

2   Method

In the example in Figure 1, three X-ray images are used. Three seeds project in each image and leave a "mark" in the image, typically a dark blotch in fluoroscopy. After reconstructing a 3D voxel volume with common back-projection, five candidate seeds are found in the volume and all appear to be legitimate seeds in each image. The question is how to separate the true seeds from the false positives. In this simplistic example, it is very easy to identify the legitimate seeds (solid circles) because any other choice will lead to an inconsistency: there will be "seed marks" in one or more X-ray images to which no seed in the volume projects. Therefore, the intuition behind the reasoning is that each "seed mark" in each image must be "covered" by at least one seed in the voxel volume. To support this intuition, we develop a theoretical framework based on optimal geometry coverage.

Tomosynthesis with Distance Map and Seed Localization. Tomosynthesis [10] is the technique to reconstruct a 3D volume from multiple 2D projection images within a limited angle. The most commonly used approach is the back projection method, which was first introduced for CT reconstruction. In this method, each voxel in a 3D volume is projected onto all the images. The value assigned to this voxel is calculated as the average of the intensity values of its projected locations in the images.

A local coordinate system must be defined for the C-arm. For any arbitrary 3D point in the space, after C-arm calibration, its projected coordinates on the 2D image plane can be calculated by the rules of perspective projection.

In practice, the C-arm's pose is estimated with some error. Since the size of the seeds is small relative to the focal length, the reconstruction is quite sensitive to pose error. To make the reconstruction more robust to pose error, in the reconstruction of the 3D voxel volume we use distance maps rather than the projection image itself.

As preprocessing, in each image we first extract the so called 2D seed regions--areas that contain seeds' projections--using adaptive thresholding and morphological operators and call the resulting images "seed-only" image. Then for each seed-only image, a distance map is calculated using a distance transform: the value at each pixel is the Euclidean distance to its nearest 2D seed region. (Pixels inside a 2D seed region all take a value of zero.) In reconstructing the 3D volume, the value of a voxel is the average of the distance values at all of its projected locations.

After the 3D voxel volume is reconstructed, candidate 3D seed regions are extracted by thresholding. The threshold value is based on estimated pose error. For example, in case of small pose error of less than 1 degree rotation, such as in Jain *et al.* [6], the threshold is set to ½ pixel. Upon thresholding, connected 3D seed regions are considered as candidate seeds. The candidate seeds are then labeled using the standard 3D "connected component labeling" method, i.e. any two neighboring voxels in the 3D seed regions are assigned the same label and considered as parts of the same seed. After that, the centroids of all the candidate seeds are calculated by averaging the 3D coordinates of all voxels wearing the same label. Each candidate seed is then represented by the location of its centroid, in addition to its label.

As we mentioned earlier, a decisive advantage of tomosynthesis over the three-film technique and its derivatives [6] is that, in addition to not requiring precise seed segmentation, it can reconstruct all hidden seeds. But the disadvantage is that tomosynthesis introduces false positive seeds, often as much as 20%. Next, we propose a theoretical framework for separating the true seeds from the false ones.

Theoretical Framework for Separating True Seeds. The problem is formulated as an optimal geometric coverage problem [11]. The general optimal coverage problem arises, for example, in wireless network design. From a given set of "server" points, we must select the minimum subset that can cover a given set of "client" points. The goal is formulated as to minimize a total cost function, which is the sum of the cost functions defined on all the selected server points.

In our problem, based on the intuition mentioned earlier, we want to find the M true seeds from the set of N candidates, such that all the 2D seed regions are "covered" in all projection images, i.e. every 2D seed region must have at least one true seed projected in it. In constructing the appropriate cost function, we can use the observation that a false positive (FP) seed, owing to its very nature, is projected close to some true seed in every image. While a true seed may be projected close to some other true seeds in some images, it usually is not the case for every image. (While it is common that a true seed can be hidden in some projection images, we have never encountered a situation where a true seed was hidden in all images. We also note that no existing method can recover such a seed.)

Hence for the server model, we define the cost function of a given seed (potential server point) as the sum of the closest distances between the projections of this seed and the projections of all other true seeds, in all images. We formulate the problem as such a general optimal coverage problem: Given the N candidate seeds, we want to find M seeds (x_1, x_2, …, x_M) from them such that all 2D seed regions in the seed-only images are covered, and the below cost function is minimized

$$C = -\sum_{m=1}^{M}\sum_{i=1}^{n}\hat{d}_i(x_m) \tag{2}$$

where
$$\hat{d}_i(x_n) = \min_m \| P^j x_m - P^j x_n \|, \qquad \forall n \in [1,2,\cdots,M], \text{and } n \neq m \tag{3}$$

and P^j is the projection operator that projects a 3D point onto the j^{th} image.

Unfortunately, the optimal coverage problem is NP-hard [12] and its computational complexity is $O(C_M^N)$, where C_M^N means N choose M. Thus for a large set of seeds, it is not possible to find its global optimal solution. We, however, managed to reduce the size of the problem by using the 2D seed-only images for regularization. To further reduce the problem, we also used greedy search to minimize local costs rather than the global cost function in Eq. (2).

Seed Clustering. To ensure that the projections of selected true seeds cover all the 2D seed regions in all seed-only images, the candidate seeds are clustered based on their projections in each image. For this purpose, we label all 2D seed regions in all seed-only images, in the similar way as 3D labeling described, e.g. the separated seed regions are assigned different labels and pixels in the same connected region wear the same label. The projections of seeds can then be clustered based on these labels. An example is shown in Fig. 2. Unlike Su *et al.* [7], the purpose of seed clustering is not to segment the 2D seed projections from 2D seed clusters in the images, or to identify the number of true seeds in each cluster. Instead, we use the seed clustering as a way to relating the 3D seeds, and the relationship is used in the coverage function that is to be minimized.
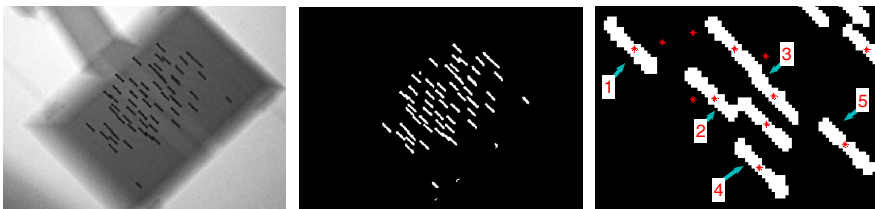

Fig. 2. (Left) X-ray image on phantom data; (Middle) the seed-only image resulted after preprocessing; (Right) Example of seed region labeling and seed grouping

On a projection image, the projections of all candidate seeds are first computed using Eq. (1). We denote a 3D candidate seed as x_n, and its 2D projection on the j^{th} image is $P^j x_n$. For each $P^j x_n$, we find its nearest seed region, which is labeled as $L^j(P^j x_n)$. The distance from $P^j x_n$ to its closest seed region is also calculated and denoted as d_n^j ($d_n^j = 0$ if $P^j x_n$ is inside a seed region). The seed projections are then

clustered based on $L^j(P^j x_n)$. Let there be K^j seed regions on the j^{th} seed-only image, and they are labeled as $1, 2, ..., K^j$. The projections of the N candidate seeds on the j^{th} image are then clustered into K^j sets Ω_k^j, such that

$$\forall p_1, p_2 \in \Omega_k^j, \ L^j(p_1) = L^j(p_1), \qquad j = 1, 2, ..., N_p, \ k = 1, 2, ..., K^j \qquad (4)$$

Let $\|\Omega_k^j\|$ be the cardinal of set Ω_k^j, we say the seed region with label k in j^{th} image is covered by $\|\Omega_k^j\|$ seeds. $\|\Omega_k^j\| \geq 1$. The clustering is repeated on all images.

In the example in Fig. 2(c), the seed regions are labeled from 1 to 5 and red asterisks mark the 2D projections of the candidate seeds. The three regions labeled with 1, 4 and 5 contain one true seed each. The two regions labeled with 2 and 3contain two (or more) true seeds each. Since the seeds in region 2 (and 3) are connected, it is considered as a single seed region. The candidate seeds include both true seeds and a few false seeds. In this example, upon clustering, the sets with label from 1 to 5 have 2, 3, 4, 1, and 1 element(s) individually

Local Coverage and Greedy Search. The seed clustering according to each projection can help reduce the size of our optimization problem. If a seed region is covered only by one seed, this seed must be a true one, because otherwise this region is not covered. The set of such seeds can be expressed as:

$$\mathbf{G} = \bigcup_j \{x_n : L^j(P^j x_n) = k \ \text{and} \|\Omega_k^j\| = 1, \ k = 1, 2, ..., K^j\}, \ \ j = 1, 2, ..., N_p \qquad (5)$$

where N_p is the number of projection images. These seeds are always chosen as true seeds. Hence the optimization problem is reduced to choose $(M-\|\mathbf{G}\|)$ true seeds from $(N-\|\mathbf{G}\|)$ candidate seeds. This can also be seen in the example in Fig. 2(c). The two regions labeled with 4 and 5 contain only one seed each. Therefore these two seeds are considered as true seeds and are always chosen.

Besides, instead of finding the global optimization of the cost function in (2), we use greedy search to find an approximate optimal solution. We also redefine the local coverage cost function as

$$C = \sum_n c(x_n), \quad c(x_n) = -\sum_{j=1}^{N_p} \frac{1 + D^j(x_n)}{1 + d_n^j}, \ \text{for all } x_n \notin \mathbf{G} \qquad (6)$$

where

$$D^j(x_n) = \min_{m \neq n} \| P^j x_n - P^j x_m \|, \quad \forall x_m \in S - G, \text{ and } L^j(P^j x_n) = L^j(P^j x_m) \qquad (7)$$

is the minimum distance between the projection of x_n and the projections of other seeds in the cluster in the j^{th} image that includes x_n, and d_n^j is the distance from $P^j x_n$ to the nearest seed region. d_n^j is added in the cost function to include the effect of imperfect X-ray pose estimation.

The minimization problem in (6) is solved using greedy search iteratively. During each cycle of iteration, the seed that has the largest cost value $c(x_n)$ is considered as a false seed and is removed from the candidate seed set. $\mathbf{G}$ is updated at the beginning of each cycle of iteration, for after the removal of one seed, there may be additional

seeds that cover some region alone (i.e. no other seed covers this region.) These seeds need to be extracted and added to **G**. The algorithm is summarized as below.

Algorithm 1. find *M* good seeds from *N* candidate seeds using greedy search
1. Initialize **S** be the set of candidate seeds.
2. For *i = 1:N-M*
3. Calculate **G** using (5).
4. Calculate local cost function $c(x)$ of all seeds *x* in **S - G** using (6).
5. Find $x_k \in \mathbf{S} - \mathbf{G}$, such that $x_k = \arg_x \min c(x)$.
6. Remove x_k from **S**: $\mathbf{S} = \mathbf{S} - \{x_k\}$.

3 Experiments and Results

Simulations. Synthetic C-arm images were used to verify our method. The images simulated a 50 cc prostate with a seed density of 2.0 seed/cc. The C-arm's focal length was 1000 mm, and the pixel size was 0.25 mm. Six images were generated on a 20^o cone around the AP axis with evenly distributed angles. The seeds were represented by cylinders with a radius of 0.4 mm and a length of 1.45 mm. A typical synthetic image is shown in Fig. 3(a). No pose estimation error was assumed.
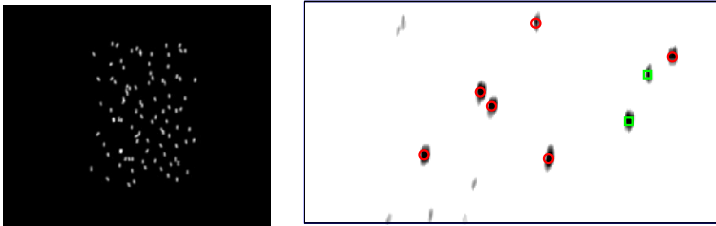

Fig. 3. (Left) Part of a synthetic C-arm image; (Right) One slice of the reconstructed 3D volume. The red circles mark the true seeds, and the green squares represent the FP seeds.

Table 1. Seed reconstruction results using different number of synthetic images

Number of images used	Number of true seeds implanted	Number of candidate seeds before FP removal	Correctly reconstructed seeds after FP removal (%)	Mean reconstruction error (mm)	Reconstruction error STD (mm)
3	96	105.6	99.4	0.19	0.19
4	96	97.7	100	0.12	0.09

We reconstructed the seeds using 3, 4 and 6 images. For 3 images, 20 experiments were performed by using all the 20 combinations of selecting 3 images from the 6 available images. For 4 images, four experiments were performed using different image combinations. The results are shown in Table 1. In each experiment, the number of implanted seeds was assumed known. After FP removal, a set of candidate seeds that equal to the number of implanted seeds were chose. These chosen seeds were then compared with the known 3D locations of implanted seeds. When using 3

images, there were about 10% FP seeds reconstructed initially and then almost all were successfully removed, indicated by the near 100% correct final reconstruction rate. This means that when there is no pose estimation error, the seeds can be reconstructed accurately even from as few as 3 images using our method.

Phantom Studies. Experiments were performed on a seed phantom constructed from acetol. The FTRAC fiducial was used to track the C-arm (with accuracy of 0.56 mm translation and 0.33^o rotation). It was attached to the seed phantom, as shown in Fig. 4. The seed phantom is comprised of twelve slabs with 5mm thickness each. Each slab has more than 100 holes with 5 mm spacing, into which the seed can be inserted. By precise manufacturing, the seed phantom was attached in a known position to a radiographic C-arm tracking fiducial, replicated after Jain *et al*. [13]. In this way, the exact location of each implanted seed was known relative to the tracking fiducial, serving as ground truth. Because some rotation error got introduced during assembly, the ground absolute seed positions had about 0.5 mm error. (Note that their relative positions were precisely known.) The seed density was about 1.56 seed/cc. Five data sets were collected with the number of seeds varying from 40 to. For each data set, six images within a 20^o cone around the AP axis were taken using a *Phillips integris V3000* fluoroscope and dewarped using the pin-cushion test. We used four, five and six images to reconstruct the seeds, by using all the combinations of the six available images, i.e. 15 experiments using four images, 6 using five images, and 1 using six images for each data set. The reconstructed seeds were compared with the computed ground truth, and the results were shown in Table 2.
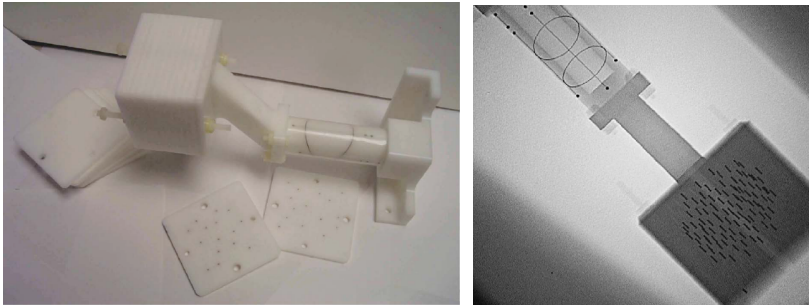

Fig. 4. (Left) An image of the seed phantom attached to the FTRAC fiducial. (Right) A typical X-ray image of the combination (with 100 seeds inserted).

It was shown in [12] that in prostate brachytherapy, 95% of the implanted seeds must be recovered in order to obtain clinically accurate estimation of the dose. Our results showed that (1) from six images more than 98% seeds can be successfully reconstructed, and (2) six images were sufficient for deriving clinically accurate dose estimation. Furthermore, the number of required images depends on the number of seeds implanted. For a smaller number of seeds (40, 55, 70), only four to five images were required. Additionally the performance of our method depends on the ratio between the number of FP seeds and the implanted seeds, which in turn depends on both the number of images used and the number of implanted seeds. Generally the lower is this ratio, the higher is the seed reconstruction ratio we may achieve.

Table 2. Seed reconstruction results on phantom data using different number of images

Number of images used	Number of seeds implanted	Number of candidate seeds before FP removal	Correctly reconstructed seeds after FP removal (%)	Mean reconstruction error (mm)	Reconstruction error STD (mm)
6	40	46	100	1.04	0.56
	55	58	100	0.67	0.39
	70	82	98.6	0.72	0.42
	85	94	100	0.97	0.44
	100	112	98	0.94	0.52
5	40	46.3	99.5	1.02	0.53
	55	68.3	98.8	0.85	0.46
	70	92.3	97.8	0.85	0.46
	85	105.3	97.0	1.00	0.57
	100	121.3	92.5	1.19	0.81
4	40	53.7	96.8	1.24	0.76
	55	82.1	94.7	0.99	0.69
	70	112	94.3	0.91	0.77
	85	135.8	90.1	1.13	0.62
	100	159.7	86.3	1.44	0.91

In summary, we presented a novel method for prostate brachytherapy seed reconstruction using C-arm images. We generate distance maps from the 2D projection images, then a 3D volume is then reconstructed using tomosynthesis using the distance maps, and finally true seeds are extracted from the voxel volume. The attractive feature of our method is that it does not require exact seed segmentation of the X-ray images. As a tradeoff, our method requires slightly higher number of images than the methods that requires elaborate explicit segmentation. Our method yields near 100% correct reconstruction from only six images with an average reconstruction accuracy of 0.86 mm (std=0.46mm). The method was robust to pose error present in radiographic C-arm tracking.

References

[1] Zelefsky, M.J., et al.: Five-year outcome of intraoperative conformal permanent I-125 interstitial implantation for patients with clinically localized prostate cancer. Int. J. Radia.t Oncol. Biol. Phy.s 67(1), 65–70 (2007)
[2] Tubic, D., et al.: Automated seed detection and three-dimensional reconstruction II., reconstruction of permanent prostate implants using simulated annealing. Med. Phys 28(11), 2272–2279 (2001)
[3] Narayanan, S., Cho, P., Marks, R.: Fast cross-projection algorithm for reconstruction of seeds in prostate brachytherapy. Med. Phys. 29(7), 1572–1579 (2002)

 [4] Todor, D., et al.: Operator-free, film-based 3D seed reconstruction in brachytherapy. Phys. Med. Biol. 47(12), 2031–2048 (2002)
 [5] Lam, S.T., et al.: Three-dimensional seed reconstruction for prostate brachytherapy using Hough trajectories. Phys. Med. Biol. 49(4), 557–569 (2004)
 [6] Jain, A., et al.: Matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL). Med. Phys. 32(11), 3475–3492 (2005)
 [7] Su, Y., et al.: Prostate brachytherapy seed localization by analysis of multiple projections: identifying and addressing the seed overlap problem. Med. Phy. 31(5), 1277–1287 (2004)
 [8] Tutar, I.B., et al.: Tomosynthesis-based localization of radioactive seeds in prostate brachytherapy. Med. Phys. 30(12), 3135–3142 (2003)
 [9] Grant, D.G.: Tomosynthesis: A three-dimensional radiographic imaging technique. IEEE Trans. Biomed. Eng. 19, 20–28 (1972)
[10] Herman, G.T., Kuba, A.: Discrete tomography in medical imaging. Proceeding of the IEEE 91(10), 1612–1626 (2003)
[11] Marengoni, M., Draper, B., Hanson, A., Sitaraman, R.: A system to place observers on a polyhedral terrain in polynomial time. Image and Vision Computing 18(10), 773–780 (2000)
[12] Su, Y., et al.: Examination of dosimetry accuracy as a function of seed detection rate in permanent prostate brachytherapy. Med. Phys. 32(9), 3049–3056 (2005)
[13] Jain, A., et al.: Robust Fluoroscope Tracking Fiducial. Med. Phys. 32(10), 3185–3198 (2005)